

## CLAIMS

The invention claimed is

1. A system comprising:

an initiator including:

a collection of files, each of the files containing data apportioned into blocks of data;

a first block level hardware interface;

a first block level communication interface configured to transmit and receive the blocks of data via the first block level hardware interface;

a storage containing data segregated into nodes; and

a snapshot server including:

a node level hardware interface communicatively linked with the storage;

a second block level hardware interface communicatively linked to the first block level hardware interface of the initiator;

a second block level communication interface configured to exchange blocks of data with the initiator via the second block level hardware interface;

a node level snapshot management configured to generate snapshots of the nodes of data contained in the storage; and

an agent associated with the collection of files, the agent configured to translate first blocks of data received from the initiator into corresponding first nodes of data to be sent via the node level hardware interface to the storage, the agent configured to translate second nodes of data received from the storage into corresponding second blocks of data to be sent to the initiator.

2. The system of claim 1 wherein the collection of files is indicated on the initiator by one of the following: a drive letter and a directory name.

3. The system of claim 1 wherein the agent is configured to convert SCSI formatted offset address associated with at least one of the first blocks of data into a terabyte/gigabyte/megabyte/sub-megabyte address of the storage associated with at least one of the first nodes of data to be sent to the storage as part of the translation of the first blocks of data into the corresponding first nodes of data.

4. The system of claim 1 wherein the initiator is a workstation.

5. The system of claim 1 wherein the first and second block level hardware interfaces are compliant with SCSI standards and the node level hardware interface is compliant with IDE standards.

6. The system of claim 1 wherein the initiator is configured to use data block sizes of 512 kilobytes.

7. The system of claim 1 wherein the agent is further configured to determine a location in the storage where to send the first nodes of data from a write command received from the initiator along with the first blocks of data.

8. The system of claim 1 wherein the agent is further configured to determine an origination location in the storage of the second nodes of data from a read command received from the initiator, and wherein the agent is configured to request from the storage transmission of the second nodes of data.

9. A system comprising:  
an initiator having a file containing data apportioned into blocks of data;  
a storage containing data segregated into nodes; and  
a snapshot server communicatively link to the initiator and the storage, the snapshot server including a node level snapshot management configured to generate

snapshots of the nodes of data contained in the storage and an agent configured to translate the blocks of data received from the initiator into corresponding nodes of data to be sent to the storage, the agent configured to translate nodes of data received from the storage into corresponding blocks of data to be sent to the initiator.

10. A system comprising:

- a first initiator configured to access a first collection of files, each file containing data apportioned into a collection of blocks of data;

- a second initiator configured to access a second collection of files, each file containing data apportioned into a collection of blocks of data;

- a storage containing data segregated into a first collection of nodes corresponding to the first collection of files and into a second collection of nodes corresponding to the second collection of files; and

- a snapshot server communicatively linked with the first initiator, the second initiator and the storage, the snapshot server including:

  - a node level snapshot management configured to generate node snapshots of the nodes of data contained in the storage, the node snapshots being stored in the storage;

  - a first agent corresponding to the first initiator configured to translate one or more of the blocks of data when received from the first initiator by the snapshot server into associated nodes of data to be sent to the storage, the first agent configured to request nodes of data from storage corresponding to a request from the first initiator for blocks of data and to translate nodes of data received from the storage into the requested blocks of data to be sent to the first initiator; and

  - a second agent corresponding to the second initiator configured to translate one or more of the blocks of data when received from the second initiator by the snapshot server into associated nodes of data to be sent to the storage, the second agent configured to request nodes of data from storage data corresponding to a request

from the second initiator for blocks of data and to translate nodes of data received from the storage into the requested blocks of data to be sent to the second initiator.

11. A system comprising:

a server;

a first block level communication;

a plurality of initiators link to the server through the first block level communication;

a second block level communication;

a storage, the server linked to the storage through the second block level communication, the storage having an initial physical storage capacity;

a plurality of agents running on the storage, each the plurality of agents configured to direct data commands from one of the plurality of initiators to the storage, each of the plurality of agents coded to indicate capacity allocation to be used to respond to queries from the plurality of initiators regarding total storage available to the querying initiator, the sum of the capacity allocations as designated being a total capacity allocation larger in size than the initial physical storage capacity of the storage; and

a monitor running on the server configured to monitor the storage for total amount of physical space on the storage being used by at least one of the plurality of initiators, the monitor configured to generate an alert when the total amount of physical space on the storage being used by at least one of the plurality of initiators reaches a predetermined fraction of the total amount of physical space available on the storage.

12. A system comprising:

an initiator having a first file indicator associated with a first collection of files and a second file indicator associated with a second collection of files, each file containing data apportioned into blocks of data;

a storage containing data segregated into nodes; and

a snapshot server communicatively linked with the initiator and the storage, the snapshot server including:

a node level snapshot management configured to generate snapshots of the nodes of data contained in the storage;

a first agent corresponding to the first file indicator configured to translate one or more blocks of data associated with the first collection of files into corresponding nodes of data to be sent to the storage, the first agent configured to request nodes of data from storage corresponding to a request from the initiator for blocks of data associated with the first collection of files and to translate nodes of data received from the storage into the requested blocks of data to be sent to the initiator; and

a second agent corresponding to the second file indicator configured to translate one or more blocks of data associated with the second collection of files into corresponding nodes of data to be sent to the storage, the second agent configured to request nodes of data from storage corresponding to a request from the initiator for blocks of data associated with the second collection of files and to translate nodes of data received from the storage into the requested blocks of data to be sent to the initiator.

13. The system of claim 12 wherein the first agent is further configured to respond to a query from the first initiator that space available on the storage is a first size, wherein the second agent is further configured to respond to a query from the second initiator that space available on the storage is a second size, the sum of the first and second size and other sizes associated with responses by other agents running on the snapshot server to queries by other initiators being larger than the physical size of the storage, the snapshot server further including a storage monitor configured to generate an alert once a predetermine amount of storage has been used.

14. A method comprising:

in a memory of an initiator, storing an indicator of a collection of files, each of the files containing data apportioned into blocks of data;

linking the initiator with a snapshot server using a first block level communication;

linking the snapshot server with a storage using a second block level communication;

in the snapshot server translating blocks of data received from the initiator into corresponding nodes of data to be sent to the storage; and

generating a snapshot of at least one of the nodes of data contained in the storage.

15. The method of claim 14 wherein storing in memory stores the indicator as one of the following: a drive letter and a directory name.

16. The method of claim 14 further including in the snapshot server, determining location in the storage where to send the nodes of data from a write command received from the initiator along with the blocks of data.

17. The method of claim 14 wherein the translating includes converting SCSI formatted offset address associated with at least one of the blocks of data into a terabyte/gigabyte/megabyte/sub-megabyte address of the storage associated with at least one of the nodes of data to be sent to the storage as part of the translation of the blocks of data into the corresponding nodes of data.

18. The method of claim 14 wherein generating a snapshot of at least one of the nodes of data includes:

copying a first node having child nodes to a new node that points to the child nodes of the first node; and

when a node is modified to generate a modified node:

replacing ancestor nodes of the modified node that have not yet been replaced with a new node;  
replacing the modified node with a new node that points to the same child nodes of the replaced node; and  
effecting the modification on the new node.

19. A method comprising:  
in memory of an initiator storing an indicator of collection of files, each of the files containing data apportioned into blocks of data;  
linking the initiator with a snapshot server using a first block level communication;  
linking the snapshot server with a storage using a second block level communication;  
in memory of the snapshot server, translating nodes of data received from the storage into corresponding blocks of data to be sent to the initiator; and  
generating a snapshot for at least one of the nodes of data contained in the storage.

20. The method of claim 17 further including in the snapshot server, determining an origination location in the storage of the nodes of data from a read command received from the initiator and from the snapshot server, requesting from storage, transmission of the nodes of data.

21. A method comprising:  
linking a first initiator having a first root directory to a snapshot server through a first block level communication;  
linking a second initiator having a second root directory to the snapshot server through a second block level communication;

linking a storage to the snapshot server through a third block level communication;

through a first agent process running on the snapshot server, pointing the first root directory to a first collection of nodes stored on the storage so that all nodes of the first collection of nodes are pointed at with a pointer from at least one of the first root directory and another of the first collection of nodes;

generating a first snapshot of the first collection of nodes;

through a second agent process running on the snapshot server, pointing the second root directory to the first snapshot so that all nodes of the first snapshot are pointed at with a pointer from at least one of the second root directory and another node of the first snapshot;

generating a second snapshot of the first snapshot;

modifying the second root directory on the second initiator to consequently modify the first snapshot.

22. The method of claim 21 further comprising:

detecting at least one node of the first snapshot having no pointer pointing at the at least one node; and

upon detection, reallocating storage space of the storage used for the at least one node as unused space.

23. A method comprising:

linking a plurality of initiators to a server through a first block level communication;

linking the server to a storage through a second block level communication, the storage having an initial physical storage capacity;

directing data commands from each of the plurality of initiators through one of a plurality of agents running on the server to the storage;



coding in each of the plurality of agents a designation indicating capacity allocation to be used to respond to queries from the plurality of initiators regarding total storage available to the querying initiator, the sum of the capacity allocations as designated being a total capacity allocation larger in size than the initial physical storage capacity of the storage;

monitoring the storage for total amount of physical space on the storage being used by at least one of the plurality of initiators; and

generating an alert when the total amount of physical space on the storage being used by at least one of the plurality of initiators reaches a predetermined fraction of the total amount of physical space available on the storage.

24. The method of claim 23 further comprising adding additional physical storage space to the storage in response to the generating an alert.